

PATENT APPLICATION

ARRAY-TYPE DISK APPARATUS PREVENTING DATA LOST WITH TWO DISK DRIVES FAILURE IN THE SAME RAID GROUP, THE PREVENTING PROGRAMMING AND SAID METHOD

Inventors: **Ikuya YAGISAWA**
Citizenship: Japan

Takeki OKAMOTO
Citizenship: Japan

Naoto MATSUNAMI
Citizenship: Japan

Mikio FUKUOKA
Citizenship: Japan

Toshio NAKANO
Citizenship: Japan

Kenichi TAKAMOTO
Citizenship: Japan

Akira YAMAMOTO
Citizenship: Japan

Assignee: **Hitachi, Ltd.**
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity: **Large**

TOWNSEND AND TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

TITLE OF THE INVENTION

ARRAY-TYPE DISK APPARATUS PREVENTING DATA LOST WITH 2
DISK DRIVES FAILURE IN THE SAME RAID GROUP, THE
PREVENTING PROGRAMMING AND SAID METHOD

5

BACKGROUND OF THE INVENTION

The present invention relates to a disk drive, which is an external memory device for a computer, and, more particularly, to a technique for preventing a plurality of disk drives in an array-type disk apparatus constituting a disk array from failing simultaneously and a technique for improving the host I/O response and improving the reliability at the time of data shifting among disk drives constituting a disk array group having a redundancy.

An array-type disk apparatus is one type of memory device systems which are to be connected to computers. The array-type disk apparatus is called a RAID (Redundant Arrays of Inexpensive Disks) and is a memory device which has a plurality of disk drives laid out in an array and a control section to control the disk drives. In the array-type disk apparatus, a read request (data read request) and a write request (data write request) are processed fast by the parallel operation of the disk drives and redundancy is added to data. As disclosed in Non-patent Publication 1 ("A Case for Redundant Arrays of Inexpensive Disks (RAID)", David A. Patterson, Garth Gibson, and Randy H. Katz,

Computer Science Division Department of Electrical
Engineering and Computer Sciences, University of
California Berkeley), array-type disk apparatuses are
classified into five levels according to the type of
5 redundant data to be added and the structure.

It is typical for array-type disk apparatuses
available on the market that spare disk drives are
mounted beforehand in the same array-type disk
apparatus on the assumption that disk drives used may
10 fail. In case where an array-type disk apparatus
decides that a disk drive which is a member of the RAID
of the array-type disk apparatus or a disk array group
has failed, the array-type disk apparatus restores the
same data and parity of the failed disk drive in the
15 associated spare disk drive based on the data and
parity of another disk drive. After restoration, the
spare disk drive operates in place of the failed disk
drive.

Further, if the data and parity of a disk drive
20 are restored after the disk drive fails, an access is
made to all the disk drives constituting the RAID group,
lowering the on-line performance. As a solution to
this problem, there is a technique which predicts a
disk drive which is likely to fail, copies data in the
25 paired spare disk drive before the disk drive fails and
becomes inaccessible, and keeps the disk operation
using the spare disk drive. Patent Document 1
(Japanese Patent Laid-Open No. 147112/1996) discloses a

technique which copies data of a disk drive to its spare disk drive and restores the data in the spare disk drive in case where the number of errors occurred in that disk drive exceeds a specified value.

5 Further, the conventional array-type disk apparatus has an operational flow such that when a data read failure occurs frequently in a disk drive from which data is shifted (hereinafter called "data-shifting disk drive") at the time of shifting data to 10 the spare disk drive of the disk drive due to preventive maintenance or so, data read from the data-shifting disk drive is attempted and after a data read failure is detected, the data in the data-shifting disk drive is restored by the disk drive that has redundancy 15 using the data restoring function of the array-type disk apparatus. It is therefore expected that the prior art drive suffers a slower response to the data read request from the host computer. To avoid the response drop, it is typical to perform the process of 20 coping with the data read request from the host computer using only the system which isolates the data-shifting disk drive from the array-type disk apparatus when a data read error has occurred frequency in the data-shifting disk drive and restores the data in the 25 data-shifting disk drive by means of the redundant disk drive by using the data restoring function of the array-type disk apparatus.

SUMMARY OF THE INVENTION

As the capacity of disk drives is ever increasing, bringing about a problem that the probability of occurrence of a data read failure in a redundant array-type disk apparatus increases in proportional to that increase. In case where a redundant array-type disk apparatus has a data unreadable portion, data in the data-shifting disk drive cannot be restored so that the data is lost as a consequence.

In case of storage system, that is, an array-type disk apparatus having redundant disk drives, i.e., one disk array group, data can be restored by using the redundancy of the array-type disk apparatus when one disk drive fails. In case of a 2 disk drives failure where, with one disk drive failing, data reading from another disk drive is disabled, data is lost.

The data restoring process of storage system, that is, an array-type disk apparatus is generally performed in parallel to an on-line process, and the capacity of the disk drives becomes larger every year, so that the data restoring time becomes longer. This increases the probability that one disk drive fails during restoration. As the capacity of the disk drives becomes larger, the time for data reading from a disk drive at the time of data restoration becomes longer, thus increasing the probability of occurrence of bit errors that cannot be recovered. It is apparent from the above that the probability of occurrence of a 2

disk drives failure of disk drives is likely to increase.

According to the prior art that copies data to a spare disk drive before its associated disk drive becomes inaccessible, if the specified value for the count of errors to be occurred which triggers the initiation of data copying to the spare disk drive is set high, the probable occurrence of possible failures is underestimated. This increases the probability of occurrence of a 2 disk drives failure. If the error count specified value level is set low, on the other hand, the frequency of usage of the spare disk drives becomes high, leading to a cost increase for the spare disk drives.

In case where an array-type disk apparatus decides that a disk drive has failed, if an attempt is made to restore the same data and parity of the failed disk drive into the spare disk drive based on the data and parity of another disk drive which is another member of the disk array group of the array-type disk apparatus but there is some data which cannot be read from that another disk drive during data restoration, data of the parity group concerning that data cannot be restored, resulting in a 2 disk drives failure.

There may be a case where while none of the disk drives constituting the disk array group of an array-type disk apparatus have not met such an event that the number of errors occurred has reached the specified

value, the numbers of errors occurred of plural disk drives approach the specified value so that it is very likely to cause a 2 disk drives failure in which some of the disk drives constituting the disk array group of the array-type disk apparatus fail at a time. The prior art that starts copying data to a spare disk drive based on the number of errors occurred cannot avoid such a possible 2 disk drives failure.

In other words, there is a case where the prior art cannot cope with a 2 disk drives failure in which some of the disk drives constituting the array-type disk apparatus fail at a time.

It is the first object of the invention to provide a highly reliable storage system, that is, a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of a 2 disk drives failure without involving a cost increase for spare disk drives.

It is the second object of the invention to provide a highly reliable array-type disk apparatus which reduces the probability of occurrence of a 2 disk drives failure when one of the disk drives constituting a disk array group has failed.

It is the third object of the invention to provide a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of

occurrence of a 2 disk drives failure when a failure potential of plural disk drives constituting the array-type disk apparatus is high.

It is the fourth object of the invention to
5 provide a highly reliable redundant array-type disk apparatus which completes data shifting without lowering the I/O response to a host computer and losing data at the time of shifting data of a disk drive in the array-type disk apparatus to its associated spare
10 disk drive.

The invention further aims at providing a control program, control method and a data shifting method which drive the array-type disk apparatuses that achieve those four objects.

15 To achieve the objects, according to the invention, there is provided an array-type disk apparatus having a plurality of disk drives, wherein at least one of the disk drives of the array-type disk apparatus is a spare disk drive, and the array-type disk apparatus has an error monitor section which monitors a status of error occurrence in each of the disk drives and instructs initiation of mirroring between that disk drive and the spare disk drive when a number of errors occurred of the disk drive exceeds a
20 specified value level 1; instructs initiation of blockade of the disk drive when the number of errors occurred of the disk drive exceeds a specified value level 2 greater than the specified value level 1, and
25

instructs shifting of a process which has been performed by the disk drive to the spare disk drive, a mirror section which performs mirroring between the disk drive and the spare disk drive, and a
5 blockade/shift section which performs blockade of the disk drive and the shifting.

The array-type disk apparatus monitors a status of error occurrence in each of the disk drives and instructs initiation of mirroring between that disk
10 drive and the spare disk drive when a number of errors occurred of the disk drive exceeds a specified value, clears mirroring of the spare disk drive when a number of errors occurred of that disk drive which is not undergoing mirroring exceeds the number of errors
15 occurred of the disk drive that is undergoing mirroring, instructs initiation of mirroring between the disk drive not undergoing mirroring and the mirroring-cleared spare disk drive, and performs mirroring between the disk drive and the spare disk drive.

20 Further, the array-type disk apparatus has an error monitor section which monitors a status of error occurrence in each of the disk drives and gives such an instruction as to set the status of the disk drive in a temporary blocked state, and a data restoring section
25 which, when a disk drive constituting a disk array group becomes the temporary blocked state, restores data of the temporary blocked disk drive from another disk drive constituting the disk array group to the

spare disk drive, and performs reading from the temporary blocked disk drive when reading from the another disk drive constituting the disk array group is not possible during data restoration.

5 Furthermore, an array-type disk apparatus having a plurality of disk drives is designed in such a way that at the time of data shifting between disk drives, a number of read errors occurred from a data-shifting disk drive is stored, data from the data-shifting disk 10 drive is read into a shifting-destination disk drives until the number of errors occurred reaches a specified value, data reading is switched to data reading from a disk drive constituting a disk array group when the number of errors occurred reaches the specified value, 15 and data reading from the data-shifting disk drive is executed when data reading from the disk drive constituting the disk array group is in error and data restoration is not possible.

20 The array-type disk apparatus monitors a status of error occurrence in each of the disk drives with a disk array group constituted by the disk drives as one unit, instructs initiation of shifting of data of that disk drive whose number of errors occurred exceeds a specified value to the spare disk drive, dynamically 25 changes the specified value to a smaller value when the numbers of errors occurred of the plurality of disk drives of the disk array group reach a sub specified value set smaller than the specified value, and

performs data copying upon reception of that shifting instruction.

The present invention can suppress the occurrence of a 2 disk drives failure in which some of the disk drives constituting a disk array (RAID) group fail at a time.

The invention has an advantage such that because the array-type disk apparatus which copies data or so to a spare disk drive for a possible failure can perform mirroring to the spare disk drive and use the spare disk drive as a spare for that disk drive which has not undergone mirroring, the probability of occurrence of a 2 disk drives failure can be reduced without involving a cost increase for spare disk drives.

The invention has another advantage such that the array-type disk apparatus which copies data or so to a spare disk drive for a possible failure can execute spontaneous switching to the spare disk drive when the number of errors occurred reaches a specified value of the second level by performing mirroring to that disk drive which has a large number of errors occurred therein from the time at which the number of errors occurred is small and dynamically changing that disk drive which is to undergo mirroring in accordance with the number of errors occurred.

The invention has a further advantage such that the probability of occurrence of a 2 disk drives failure can be reduced in a disk array system in which

one of disk drives constituting a disk array (RAID) group fails.

5 The invention has a still further advantage such that the probability of occurrence of a 2 disk drives failure can be reduced in an array-type disk apparatus which copies data or so to a spare disk drive for a possible failure in a state where the failure potential of plural disk drives constituting the array-type disk apparatus is high.

10 Furthermore, the invention has an advantage such that at the time of shifting data among disk drives in a large-capacity array-type disk apparatus, the hybrid system of a data restoring system based on redundant data and system of reading from a data-shifting disk 15 drive can shift data to the data-shifting disk drive method without losing it by keeping using the data-shifting disk drive without being completely isolated.

BRIEF DESCRIPTION OF THE DRAWINGS

20 Fig. 1 is a structural diagram of an array-type disk apparatus according to a first embodiment of the invention;

25 Fig. 2 is an explanatory diagram of a disk drive management table according to the first embodiment of the invention;

Fig. 3 is an explanatory diagram of disk drive management means according to the first embodiment of the invention;

Fig. 4 is a flowchart of a preventive spare copying operation according to the first embodiment of the invention;

5 Fig. 5 is an explanatory diagram of a disk drive management table according to a second embodiment of the invention;

Fig. 6-1 is a flowchart of a dynamic mirroring operation according to the second embodiment of the invention;

10 Fig. 6-2 is a flowchart of the dynamic mirroring operation according to the second embodiment of the invention;

15 Fig. 7 is a structural diagram of an array-type disk apparatus according to a third embodiment of the invention;

Fig. 8 is an explanatory diagram of a disk drive management table according to the third embodiment of the invention;

20 Fig. 9 is an explanatory diagram of a disk drive management section according to the third embodiment of the invention;

Fig. 10 is a flowchart of a sector failure restoring operation according to the third embodiment of the invention;

25 Fig. 11 is a flowchart of a write operation in the sector failure restoring operation according to the third embodiment of the invention;

Fig. 12 is an explanatory diagram of a disk drive

management table according to a fourth embodiment of the invention;

Fig. 13 is an explanatory diagram of disk drive management means according to the fourth embodiment of the invention;

Fig. 14 is a flowchart of a 2 disk drives failure preventing operation according to the fourth embodiment of the invention;

Fig. 15 is a diagram showing the drive structure according to a fifth embodiment of the invention;

Fig. 16 is a diagram showing the details of the drive structure according to the fifth embodiment of the invention;

Fig. 17 is a diagram showing the details of a part of the drive structure according to the fifth embodiment of the invention;

Fig. 18 is an operational flowchart according to the fifth embodiment of the invention;

Fig. 19 is another operational flowchart according to the fifth embodiment of the invention;

Fig. 20 is a different operational flowchart according to the fifth embodiment of the invention; and

Fig. 21 is a flowchart illustrating a further principle of the invention according to the fifth embodiment of the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS (First Embodiment)

The first embodiment of the invention is designed to achieve the first object of the invention.

That is, the first embodiment aims at providing highly reliable storage system, that is, a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of a 2 disk drives failure without involving a cost increase for spare disk drives.

10 (1) Description of Structure

The system structure of the first embodiment of the invention is discussed below referring to Figs. 1 to 3. In Fig. 1, "100" denotes a host computer, "123" denotes an array-type disk apparatus, "200" denotes the management control section of the array-type disk apparatus, "310" denotes a group of disk drives and "500" denotes a management console.

20 The array-type disk apparatus 123, the host computer 100, the management control section 200, the disk drive group 310 and the management console 500 are connected to one another in the illustrated manner.

The array-type disk apparatus 123 includes the following components as the management control section 200. The management control section 200 includes a CPU 201 which controls the management control section 200, a memory 202, a cache 203 which buffers data of a user, a host interface (I/F) 204 which executes data transmission and reception with respect to the host

computer 100, a disk drive I/F 205 which executes data transmission and reception with respect to the disk drive group 310, and a management I/F 207 which executes transmission and reception of control with respect to the management console 500. Those components are connected to one another as illustrated. The memory 202 has a RAID control section 210 which controls the disk array, a disk drive management section 230 which manages the disk drive group 310, a disk drive management table 240 which records disk drive information such as the operational parameters and operation statuses of the disk drive group 310, a disk drive information setting section 250 which sets disk drive information upon reception of an input from the management console 500, and a disk drive information notifying section 260 notifies disk drive information as an output to the management console 500.

The disk drive group 310 comprises disk drives 301 to 307. The disk drives 301 to 305 constitutes a disk array group which has the performance and reliability enhanced by the parallel operation and redundancy of disks that have been discussed in the foregoing description of the embodiment and this state is said to be constructing a disk array group to be a RAID group with the set of the disk drives 301 to 305. The disk drives 306 and 307 are spare disk drives that are placed in the disk array group in place of those disk drives constituting the disk array (RAID) group

which fail.

The management console 500 comprises an input section 510 which inputs user's settings to the disk drives 301 to 305 and an output section 520 which 5 informs the user of the information of the disk drives 301 to 305. Disk drive operation parameters to the disk drive management table 240 are input from the input section 510. The output section 520 outputs and displays the disk drive operational statuses of the 10 disk drive management table 240.

Fig. 2 shows the disk drive management table 240. The parameters include "disk drive No." which represents the identification (ID) number of each disk drive, "error counter" which stores the accumulated 15 number of errors of a disk drive, "error count specified value level 1" indicating the value of the first level as the index for the accumulated number of errors of a disk drive, "error count specified value level 2" indicating the value of the second level as 20 the index for the accumulated number of errors of a disk drive, "spare bit" indicating usage as a spare disk drive, "disk drive status" indicating the operational status of a disk drive, and "pair disk drive" indicating the association with a spare disk 25 drive which is used to cope with a disk drive failure.

Set in the "error count specified value level 1" is a value indicating the timing to start mirroring with the spare disk drive when the number of errors of

a target disk drive is accumulated and it becomes very likely to cause a failure. Set in the "error count specified value level 2" is a value which is higher than the value of the "error count specified value level 1", a value indicating the timing to block the disk drive and end mirroring with the spare disk drive as a value for determining that the number of errors of a target disk drive is accumulated and the continuous operation does not seem possible. "YES" is set in the "spare bit" when the disk drive in question is the spare disk drive and "NO" is set otherwise. The "error count specified value level 1", "error count specified value level 2" and "spare bit" are set by the user using the input section 510 of the management console 500.

Set in the "disk drive status" are a parameter "normal" indicating that the operational status of a disk drive is not abnormal, a parameter "mirror" indicating that mirroring with the spare disk drive is being done, and a parameter "blocked" indicating that the value of the error counter has reached the "error count specified value level 2" and the continuous operation of the disk drive does not seem possible. The "disk drive No." of the disk drive which becomes a pair in mirroring is set in the "pair disk drive". The individual parameter values of the disk drive management table 240 are output and displayed on the output section 520 of the management console 500 in

response to an instruction from the user.

Fig. 3 shows the disk drive management section 230. An error monitor section 231 monitors the status of occurrence of errors of a disk drive, instructs 5 initiation of mirroring of the disk drive with the spare disk drive when the number of errors occurred of the disk drive exceeds the "error count specified value level 1", and instructs termination of mirroring when the number of errors exceeds the "error count specified 10 value level 2". An error counter section 232 counts the number of errors occurred of the disk drive and sets the counted number of errors occurred to the "error counter" in the disk drive management table 240. An error-count specified value setting section 233 sets 15 a parameter, designated by the user using the management console 500, to the disk drive management table 240. A disk drive status setting section 234 sets the operational status of a disk drive to the disk drive management table 240 in response to an 20 instruction from the error monitor section 231. A mirror section 235 performs mirroring of an access from one disk drive to the spare disk drive. A blockade/shift monitor section 236 instructs blockade of a disk drive and shifting of the process which is 25 being performed by the disk drive to the spare disk drive. A blockade/shift section 237 performs blockade and shifting of a disk drive in response to an instruction from the blockade/shift monitor section 236.

The above has discussed the system structure of the array-type disk apparatus according to the embodiment.

(2) Preventive Spare Copying Operation

5 The prior art drive monitors the number of errors occurred of a disk drive, copies data of that disk drive to a spare disk drive when the number of errors reaches a certain specified value and blocks the disk drive, whereas the first embodiment has two levels of
10 specified values and starts mirroring with the spare disk drive when the number of errors occurred reaches the first specified value level 1. At this time, the disk drive is not blocked but kept operating. When the number of errors occurred reaches the second specified
15 value level 2, mirroring is cleared, the disk drive is blocked and the operation continues with the spare disk drive.

The preventing spare copying operation is discussed below using a flowchart in Fig. 4.

20 It is premised on that the error occurrence statuses of the individual disk drives 301 to 307 are counted by the error counter section 232 and are continuously set in the disk drive management table 240. The flowchart in Fig. 4 should be executed
25 independently for the disk drives 301 to 305 constituting the disk array (RAID) group.

First, the error monitor section 231 determines whether or not the value of the "error counter" in the

disk drive management table 240 of that disk drive which is to be monitored (hereinafter also referred to as "target disk drive") has reached the "error count specified value level 1" (step 1001). When the former value has not reached the error count specified value level 1, step 1001 is repeated. When the former value has reached the error count specified value level 1, a disk drive whose "spare bit" is YES is searched for and a spare disk drive is selected (step 1002). Thereafter, 5 the error monitor section 231 sets the disk drive number of the target disk drive in the "pair disk drive" of the selected spare disk drive (step 1003), and sets the number of the spare disk drive into the "pair disk drive" of the target disk drive (step 1004). 10 Next, the error monitor section 231 sets the "disk drive status" of the target disk drive and the spare disk drive in the mirror status (step 1005), and 15 instructs the mirror section 235 to start mirroring of the target disk drive and the spare disk drive (step 1006). 20

Fig. 2 shows an example of the settings in the disk drive management table 240. In the disk array in which the disk array (RAID) group is comprised of disk drives having "disk drive Nos." 0 to 4, the disk drive with the "disk drive No." 4 has an "error counter" 25 value of 60 exceeding the value "50" which is the "error count specified value level 1". This is the state where mirroring with the disk drive with the

"disk drive No." 5 or a spare disk drive has already started, the "disk drive status" of the disk drive with the "disk drive No." 4 is "mirror" and its "pair disk drive" is the disk drive with the "disk drive No." 5, 5 while the "disk drive status" of the disk drive with the "disk drive No." 5 is "mirror" and its "pair disk drive" is the disk drive with the "disk drive No." 4.

Returning to Fig. 4, in the next step, the error monitor section 231 determines whether or not the value 10 of the "error counter" in the disk drive management table 240 of the target disk drive has reached the "error count specified value level 2" (step 1007). When the former value has not reached the error count specified value level 2, step 1007 is repeated. When 15 the former value has reached the error count specified value level 2, the blockade/shift monitor section 236 instructs initiation of blockade and initiation of shifting to the spare disk drive, and sets the "disk drive status" of the target disk drive to the blocked status and the "disk drive status" of the spare disk drive to the normal status (step 1008), then instructs the mirror section 235 to terminate mirroring of the target disk drive and the spare disk drive and shifts the process which has been executed by the target disk 20 drive to the spare disk drive (step 1009). The blockade and shifting are carried out by the blockade/shift section 237. To check from which disk 25 the shifting to the spare disk is done, the value of

the "pair disk drive" should be referred to.

The above has explained the preventing spare copying operation.

(3) Advantages

5 The prior art drive monitors the number of errors occurred of a disk drive, copies data of that disk drive to a spare disk drive when the number of errors reaches a certain specified value and blocks the disk drive, whereas the first embodiment has two levels of
10 specified values and starts mirroring with the spare disk drive when the number of errors occurred reaches the first specified value level. At this time, the disk drive is not blocked but kept operating. When the number of errors occurred reaches the second specified
15 value level, mirroring is cleared, the disk drive is blocked and the operation continues with the spare disk drive.

Because the target disk drive and the spare disk drive merely undergo mirroring, if a disk drive other
20 than the target disk drive has an error occurrence status exceeding the second specified value level, it is possible to clear mirroring of this target disk drive and use the spare disk drive as a spare for another disk drive.

25 It is assumed that as shown in the example of the settings in the disk drive management table 240 in Fig. 2, for example, the disk drive with the "disk drive No." 4 has an "error counter" value of 60 exceeding the

value "50" which is the "error count specified value level 1" and the disk drive with the "disk drive No." 4 and the disk drive with the "disk drive No." 5 are subjected to mirroring. In this state, in case where 5 the value of the "error counter" of the disk drive with the "disk drive No." 0 exceeds the value of "90" which is the "error count specified value level 2", the error monitor section 231 can clear mirroring with the disk drives with the "disk drive Nos." 4 and 5 and can use 10 the disk drive with the "disk drive No." 5 as a spare for the disk drive with the "disk drive No." 0. As the frequency of occurrence of errors of the disk drive with the "disk drive No." 0 becomes higher so that the disk drive is likely to fail, data is copied to the 15 spare disk drive before the disk drive actually fails.

Because a spare disk drive can be used as a spare for another disk drive, the first specified value level can be set lower than the value specified in the prior art and the resistance to the 2 disk drives failure can 20 be improved. As the spare disk drive can be used for different disk drives, the cost for the spare disk drives can be suppressed as compared with the prior art which blocks the target disk drive in the first level.

As mirroring is performed in the first level, it 25 is possible to spontaneously switch to the spare disk drive when the number of errors reaches the second specified value level.

In short, the first embodiment can provide a

highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of a 2 disk drives failure without involving a cost increase for
5 spare disk drives.

(Second Embodiment)

The second embodiment, like the first embodiment, is designed to achieve the first object of the invention. That is, the second embodiment aims at
10 providing highly reliable storage system, that is, a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of a 2 disk drives failure without involving a cost increase for
15 spare disk drives.

(1) Description of Structure

The system structure of the second embodiment of the invention is discussed below. For the sake of descriptive simplicity, only the differences from the
20 first embodiment are discussed below. The system structure is the same as that of the first embodiment as shown in Fig. 1.

The disk drive group 310 comprises disk drives 301 to 307. The disk drives 301 to 305 constitutes a
25 disk array whose performance and reliability are enhanced by the parallel operation and redundancy of disks that have been discussed in the foregoing description of the embodiment and this state is said to

be constructing a disk array group to be a RAID group with the set of the disk drives 301 to 305. The disk drives 306 and 307 are spare disk drives that are placed in the disk array (RAID) group in place of those 5 disk drives constituting the disk array (RAID) group which fail. The second embodiment differs from the first embodiment in that mirroring is performed on that disk drive which has a large number of errors occurred from a point of time at which the number of errors 10 occurred was small. While it is desirable that all the spare disk drives or two or more spare disk drives should be subjected to mirroring, a single spare disk drive will do. In case where the number of errors occurred in a disk drive other than those disk drives 15 which are being mirrored exceeds the numbers of errors occurred in the mirroring disk drives, mirroring of that mirroring disk drive which has the smallest number of errors occurred is cleared and the mirroring-cleared disk drive is used as a spare disk drive for mirroring 20 of the disk drive whose number of errors occurred becomes large. As a disk drive to be mirrored is dynamically switched, this operation is called "dynamic mirroring operation".

Fig. 5 shows the disk drive management table 240 25 of the second embodiment, and the parameters are the same as those of the first embodiment shown in Fig. 2. The second embodiment differs from the first embodiment in that set in the "error count specified value level

1" is a value indicating the timing to check the "error counters" of all the disk drives and start mirroring of the spare disk drive and that disk drive which has a higher "error counter" when the number of errors 5 occurred in a target disk drive is accumulated and the probability of occurrence of a 2 disk drives failure of the disk drive becomes high.

Set in the "disk drive status" are a parameter "normal" indicating that the operational status of a 10 disk drive is not abnormal, a parameter "mirror" indicating that mirroring with the spare disk drive is underway, and a parameter "blocked" indicating that the value of the error counter has reached the "error count specified value level 2" and the continuous operation 15 of the disk drive does not seem possible.

In the second embodiment, the disk drive management section 230 is as illustrated in Fig. 3, and the error monitor section 231 monitors the status of occurrence of errors of a disk drive, and checks the 20 "error counters" of all the disk drives and starts mirroring of the spare disk drive and that disk drive which has a higher "error counter" when the number of errors occurred in a target disk drive exceeds the "error count specified value level 1", and instructs 25 termination of mirroring when the number of errors occurred exceeds the "error count specified value level 2".

The above is the description of the system

structure of the embodiment.

(2) Dynamic Mirroring Operation

The prior art drive monitors the number of errors occurred of a disk drive, copies (mirrors) data of that disk drive to a spare disk drive when the number of errors reaches a certain specified value and blocks the disk drive, whereas the second embodiment performs mirroring on that disk drive which has a large number of errors occurred from a point of time at which the number of errors occurred was small, and dynamically switches a disk drive to be mirrored in accordance with the number of errors occurred.

The dynamic mirroring operation is described next using flowcharts in Figs. 6-1 and 6-2. It is premised on that the error occurrence statuses of the individual disk drives 301 to 307 are counted by the error counter section 232 and are continuously set in the disk drive management table 240.

First, the error monitor section 231 determines whether or not there is a disk drive the value of whose "error counter" in the disk drive management table 240 has reached the "error count specified value level 1" (step 1501). In this case, it does not matter which disk drive has the "error counter" value that has reached the "error count specified value level 1". In case where there is no disk drive whose "error counter" value has reached the "error count specified value level 1", step 1501 is repeated.

In case where there is a disk drive whose "error counter" value has reached the "error count specified value level 1", the values of the "error counter" of all the disk drives are checked (step 1502). Next, the 5 error monitor section 231 searches for a disk drive whose "spear bit" is "YES" and determines whether or not there is any disk drive whose "mirror status" is "mirror", i.e., an unpaired spare disk drive (step 1503).

10 When there is an unpaired spare disk drive, the error monitor section 231 selects that one of unpaired disk drives whose "error counter" value is the largest as a paring target (step 1504), sets the number of the target disk drive in the "pair disk drive" of the 15 selected spare disk drive (step 1505), sets the number of the spare disk drive into the "pair disk drive" of the target disk drive (step 1506), sets the "disk drive statuses" of the target disk drive and the spare disk drive in the mirror status (step 1507), instructs the mirror section 235 to start mirroring (step 1508), then 20 returns to step 1503.

When there is no unpaired spare disk drive, the flow goes to step 1509.

Fig. 5 shows an example of the settings in the 25 disk drive management table 240. In the disk array in which the RAID group is comprised of disk drives having "disk drive Nos." 0 to 4, the disk drive with the "disk drive No." 2 has an "error counter" value of 35

exceeding the value "30" which is the "error count specified value level 1". This is the state where the flow has already proceeded to step 1509 and mirroring of the disk drive with the "disk drive No." 5 or a
5 spare disk drive and a disk drive with the "disk drive No." 2 has already started, the "disk drive status" of the disk drive with the "disk drive No." 2 is "mirror" and its "pair disk drive" is the disk drive with the "disk drive No." 5, while the "disk drive status" of the disk drive with the "disk drive No." 5 is "mirror" and its "pair disk drive" is the disk drive with the
10 "disk drive No." 2. It is also the state where mirroring of the disk drive with the "disk drive No." 4 which has the second largest "error counter" value and a disk drive with the "disk drive No." 6 or a spare disk drive has already started, the "disk drive status" of the disk drive with the "disk drive No." 4 is
15 "mirror" and its "pair disk drive" is the disk drive with the "disk drive No." 6, while the "disk drive status" of the disk drive with the "disk drive No." 6 is "mirror" and its "pair disk drive" is the disk drive with the "disk drive No." 4.
20

25 Returning to Fig. 6-2, as the next step 1509, the error monitor section 231 determines whether or not a disk drive whose "error counter" value exceeds that of the paired disk drive is included in unpaired disk drives (step 1509).

When there is such a disk drive, the error

monitor section 231 selects that one of unpaired disk drives whose "error counter" value exceeds that of a paired disk drive as a paring target (step 1510),
5 clears pairing of that of the paired disk drives whose "error counter" value is the smallest (step 1511), sets the number of the target disk drive in the "pair disk drive" of the pairing-cleared spare disk drive (step 1512), sets the number of the spare disk drive into the "pair disk drive" of the target disk drive (step 1513),
10 sets the "disk drive statuses" of the target disk drive and the spare disk drive in the mirror status (step 1514), instructs the mirror section 235 to start mirroring (step 1515), then returns to step 1509.

15 Steps 1509 to 1515 are explained below using an example of the settings in the disk drive management table 240 shown in Fig. 5. This diagram shows the state where mirroring of the disk drive with the "disk drive No." 5 which is a spare disk drive and the disk drive with the "disk drive No." 2 is carried out and
20 mirroring of the disk drive with the "disk drive No." 6 which is a spare disk drive and the disk drive with the "disk drive No." 4 is carried out.

25 It is assumed that under the situation, the "error counter" value of the disk drive with the "disk drive No." 0 is 25 which exceeds that of any mirrored disk drive. In this case, the decision in step 1509 is YES, the next mirroring target is the disk drive with the "disk drive No." 0, pairing of the disk drive with

the "disk drive No." 4 or that one of the mirrored disk drives whose "error counter" value is the smallest is cleared, and mirroring of the disk drive with the "disk drive No." 6 which is the pairing-cleared spare disk 5 drive and the disk drive with the "disk drive No." 0 is executed.

Returning to Fig. 6-2, when a disk drive whose "error counter" value exceeds that of the paired disk drive is not included in the unpaired disk drives in 10 step 1509, the error monitor section 231 determines whether or not the value of the "error counter" of the target disk drive has reached the "error count specified value level 2" (step 1516). When the former value has not reached the error count specified value 15 level 2, the flow returns to step 1509. When the former value has reached the error count specified value level 2, the "disk drive status" of the target disk drive is set to the blocked status and the "disk drive status" of the spare disk drive is set to the normal status (step 1517), an instruction is sent to 20 the mirror section 235 to terminate mirroring of the target disk drive and the spare disk drive and the process which has been executed by the target disk drive is shifted to the spare disk drive (step 1518), 25 after which the flow returns to step 1509. To check from which disk the shifting to the spare disk is done, the value of the "pair disk drive" should be referred to.

The dynamic mirroring operation is performed as described above.

With the value of the "error count specified value level 1" set to 0, the dynamic mirroring 5 operation starting at step 1502 may be executed from the beginning. The criterion for the decision in step 1509 may be the determination of whether or not a disk drive whose "error counter" value exceeds the maximum value of the "error counters" of the paired disk drives 10 is included in unpaired disk drives. Alternatively, the step 1509 may determine whether or not a disk drive whose "error counter" value exceeds an intermediate value, an average value or so derived from the "error counter" values of the paired disk drives is included 15 in unpaired disk drives.

(3) Advantages

The prior art drive monitors the number of errors occurred of a disk drive, copies data of that disk drive to a spare disk drive when the number of errors 20 reaches a certain specified value and blocks the disk drive, whereas the second embodiment executes mirroring on that disk drive which has a large number of errors occurred from a point of time at which the number of errors occurred was small and dynamically switches the 25 disk drive in accordance with the number of errors occurred. This increases the probability of instantaneously switching to a spare disk drive when the number of errors occurred reaches the second

specified value level and thus improves the resistance to a 2 disk drives failure of disk drives.

Although the foregoing description has described that dynamic mirroring is performed with respect to a single disk array (RAID) group, dynamic mirroring may be performed with respect to the entire disk array (RAID) group in the array-type disk apparatus using all the spare disk drives in the array-type disk apparatus.

(Third Embodiment)

10 The third embodiment is designed to achieve the second object of the invention.

That is, the third embodiment aims at providing highly reliable storage system, that is, a highly reliable array-type disk apparatus which reduces the probability of occurrence of a 2 disk drives failure when one of the disk drives constituting a disk array (RAID) group has failed.

(1) Description of Structure

20 The system structure of the third embodiment of the invention is discussed below using Figs. 7 to 9. For the sake of descriptive simplicity, only the differences from the first embodiment are discussed below. In Fig. 7, a data restoring section 270 which, when a disk drive is blocked, restores data from another disk drive constituting a disk array (RAID) group to a spare disk drive is provided in the memory 202 in addition to the structure in Fig. 1.

The parameters in the disk drive management table

240 in Fig. 8 are the parameters in Fig. 2 from which the error count specified value level 2 is omitted. The contents of the parameters in Fig. 8 differ from those in Fig. 2 in the following points.

- 5 Set in the "error count specified value level 1" is a value indicating the timing to start copying to the spare disk drive when the number of errors occurred in a target disk drive is accumulated and the possibility of occurrence of a failure becomes high.
- 10 After copying ends, the processing of the target disk drive is shifted to the spare disk drive but reading from the target disk drive which is carried out by the data restoring section 270 is permitted.
- 15 Set in the "disk drive status" are a parameter "normal" indicating that the operational status of a disk drive is not abnormal, a parameter "copy" indicating that the error counter value has reached the "error count specified value level 1" and copying to the spare disk drive is underway, a parameter
- 20 "temporary blocked" indicating that copying to the spare disk drive has finished and reading from the target disk drive which is carried out by the data restoring section 270 is permitted, a parameter "blocked" indicating that copying is finished, and a parameter "restoring" indicating that a process of restoring data from another disk drive constituting the disk array (RAID) group to the spare disk drive is underway. A parameter "disk drive No." of a disk drive
- 25

to be a pair to which copying is to be done is set in the "pair disk drive".

Fig. 9 shows the disk drive management section 230 according to the third embodiment and has a copy section 238 in place of the mirror section 235 in Fig.

3. The error monitor section 231 monitors the status of occurrence of errors of a disk drive, instructs initiation of copying to a spare disk drive from a target disk drive when the number of errors occurred in the target disk drive exceeds the "error count specified value level 1", sets the "temporary blocked" status during copying and sets the "blocked" status after copying is done. The copy section 238 copies data in one disk drive to a spare disk drive.

15 The above is the description of the system structure of the embodiment.

(2) Sector Failure Restoring Operation

This embodiment improves the data restoration capability in case of a 2 disk drives failure where with one sector becoming unreadable so that data is to be restored to a spare disk drive from another disk drive constituting the disk array (RAID) group, one sector in said another disk drive constituting the disk array (RAID) group further becomes unreadable. The disk drive one sector of which has become unreadable is set to the "temporary blocked" status where reading executed by the data restoring section 270 is permitted.

The sector failure restoring operation is

discussed next using a flowchart in Fig. 10. It is premised on that the error occurrence statuses of the individual disk drives 301 to 307 are counted by the error counter section 232 and are continuously set in 5 the disk drive management table 240. The flowchart in Fig. 10 should be executed independently for the disk drives 301 to 305 constituting a disk array group. The disk drive with the "disk drive No." 4 constituting the disk array (RAID) group has its number of errors 10 increasing and has one sector having become unreadable and is thus set to the "temporarily blocked" status regardless of the error counter. It is assumed that data is being restored to the spare disk drive with the "disk drive No." 5 using the disk drives with the "disk 15 drive Nos." 0 to 3 and the redundancy of the disk array (RAID). It is further assumed that under the situation, one sector of the disk drive with the "disk drive No." 0 becomes unreadable so that data is read from the same sector in the disk drive with the "disk drive No." 4 to 20 restore the disk array (RAID) group.

First, based on the data of the disk drive with the "disk drive Nos." 0 to 3, the data restoring section 270 starts a data restoring process, equivalent to a data restoring process to be done on the disk 25 drive with the "disk drive No." 4, with respect to the spare disk drive with the "disk drive No." 5 (step 2001). Next, the data restoring section 270 determines whether or not restoration is finished (step 2002).

When the restoration is finished, the data restoring section 270 shifts the processing of the disk drive with the "disk drive No." 4, which is the restoration target, to the spare disk drive (step 2003), then 5 terminates the process (step 2004). When the restoration has not ended, the data restoring section 270 determines whether or not the disk drives with the "disk drive Nos." 0 to 3 have a sector failure which disables sector reading (step 2005). When there is no 10 sector failure, step 2002 is repeated. When there is a sector failure, the data restoring section 270 attempts to read data from the same sector in the disk drive with the "disk drive No." 4 which is in the "temporary blocked" status (step 2006). The data restoring 15 section 270 determines whether or not reading is successful (step 2007), and executes a restoring process based on the contents of the read sector (step 2008) and returns to step 2002 when reading is successful. When reading is failed, the corresponding 20 sector is treated as data lost (step 2009) after which the flow returns to step 2002.

The sector failure restoring operation is performed as described above.

(3) Write Operation in Sector Failure Restoring 25 Operation

Suppose, as the premise, that the error occurrence statuses of the individual disk drives 301 to 307 are counted by the error counter section 232 and

are continuously set in the disk drive management table 240. It is assumed that the flowchart in Fig. 11 is performed on the entire disk array (RAID) group comprised of the disk drives 301 to 305. Further, the 5 disk drives 301 to 305 constitutes the disk array (RAID) group, data and a parity are stored in each disk drive and a set of a parity and data for computing the parity is called "stripe set".

Referring to Fig. 11, when the management control 10 section 200 receives a write request from the host computer 100, the disk array (RAID) control section 210 determines whether or not a writing destination is a temporary blocked disk drive (step 2501).

When the writing destination is a temporary 15 blocked disk drive, the processes starting at step 2502 take place. Suppose that the disk drive 305 is the temporary blocked disk drive and the disk drive 301 is the disk drive where the parity in the same stripe set as having the data to be written (or write data) is stored. First, the RAID control section 210 reads data 20 in the same stripe set corresponding to the write data from the disk drives 302 to 304 other than the temporary blocked disk drive 305 and the disk drive 301 where the parity is stored (step 2502). Next, the 25 exclusive-OR of the write data and the data read in step 2502 is computed, thus generating a new parity (step 2503). Then, the write data is written in the disk drive 305 or the temporary blocked disk drive

(step 2504), and the new parity is stored in the parity-stored disk drive 301 (step 2505) after which the processing is terminated.

When the writing destination is not a temporary
5 blocked disk drive, the processes starting at step 2507
take place. The RAID control section 210 determines
whether or not the parity in the same stripe set as
having the write data is located in the temporary
blocked disk drive (step 2507).

When the parity is located in the temporary
10 blocked disk drive, the processes starting at step 2508
take place. Suppose that the disk drive 305 is the
temporary blocked disk drive and the disk drive 301 is
the disk drive where write data is stored. First, the
15 RAID control section 210 reads data in the same stripe
set corresponding to the write data from the disk
drives 302 to 304 other than the temporary blocked disk
drive 305 and the disk drive 301 where data is stored
(step 2508). Next, the exclusive-OR of the write data
20 and the data in the same stripe set read in step 2508
is computed, thus generating a new parity (step 2509).
Then, the write data is written in the disk drive 301
25 (step 2510), and the new parity is stored in the disk
drive 305 which is the parity-stored disk drive where
the parity is stored (step 2511) after which the
processing is terminated.

When the parity is not located in the temporary
blocked disk drive, the processes starting at step 2512

take place. Suppose that the disk drive 305 is the temporary blocked disk drive, the disk drive 301 is the disk drive where write data is stored, and the disk drive 302 is the disk drive where the parity in the same stripe set is stored. First, the RAID control section 210 reads old data from the disk drive 301 where write data before update is stored and reads an old parity from the disk drive where a parity before update is stored (step 2512). Next, the exclusive-OR 5 of the write data, the old data and the old parity, the latter two read in step 2512, is computed, thus generating a new parity (step 2513). Then, the write data is written in the disk drive 301 (step 2514), and the new parity is stored in the disk drive 302 where 10 the parity is stored (step 2515) after which the processing is terminated.

15

The above is the description of the write operation when a write request is issued from the host computer 100 during sector failure restoration.

20 As data can be restored by using the redundancy of the disk array (RAID), writing to a temporarily blocked disk drive in step 2504 and step 2511 may be omitted. Instead of writing to a temporarily blocked disk drive in step 2504 and step 2511, writing may be 25 done to the spare disk drive to which spare copying is being performed. In addition to writing to a temporarily blocked disk drive in step 2504 and step 2511, the contents of the temporarily blocked disk

drive may be written into the spare disk drive which is undergoing spare copying.

(4) Advantages

The third embodiment can improve the data 5 restoration capability in case of a 2 disk drives failure where with one sector becoming unreadable so that data is to be restored to a spare disk drive from another disk drive constituting the disk array (RAID) group, one sector in said another disk drive 10 constituting the disk array (RAID) group further becomes unreadable.

In short, the embodiment can provide a highly reliable array-type disk apparatus which reduces the probability of occurrence of a 2 disk drives failure 15 when one of the disk drives constituting a disk array (RAID) group has failed.

Although the foregoing description has been given on the premise that preventive copying is performed on a spare disk drive, this embodiment can be adapted to 20 an array-type disk apparatus which does not perform preventive copying.

Although restoration in an array-type disk apparatus in the foregoing description is initiated on the premise that one sector of a disk drive becomes 25 unreadable, other conditions may be employed. For example, restoration of a disk drive may be initiated when that disk drive is considered as being in a blocked status as the number of errors occurred in the

disk drive has exceeded the specified value.

(Fourth Embodiment)

The fourth embodiment is designed to achieve the third object of the invention.

5 That is, the fourth embodiment aims at providing highly reliable storage system, that is, a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of 2 disk drives 10 failure when a failure potential of plural disk drives constituting a disk array (RAID) group is high.

(1) Description of Structure

The system structure of the fourth embodiment of the invention is discussed below using Figs. 12 and 13.

15 For the sake of descriptive simplicity, only the differences from the first embodiment are discussed below. The structure of this array-type disk apparatus is the same as that of the second embodiment in Fig. 7, except that the data restoring section 270 need not 20 have a function of reading a sector associated with a sector failure when the failure occurs during data restoration.

The parameters in the disk drive management table 240 in Fig. 12 are the parameters in Fig. 8 to which an 25 error count sub specified value is added. The contents of the parameters in Fig. 12 differ from those in Fig. 8 in the following points.

Set in the "error count specified value level 1"

is a value indicating the timing to start copying to the spare disk drive when the number of errors occurred in a target disk drive is accumulated and the possibility of occurrence of a failure becomes high.

5 After copying ends, the processing of the target disk drive is shifted to the spare disk drive and the target disk drive is set to a blocked status. The "error count sub specified value" is set to a value lower than the "error count specified value level 1" and when the 10 numbers of errors occurred in plural disk drives in those disk drives constituting a disk array (RAID) group reach the error count sub specified value, it means that those disk drives are potentially very likely to fail at the same time.

15 Set in the "disk drive status" are a parameter "normal" indicating that the operational status of a disk drive is not abnormal, a parameter "copy" indicating that the error counter value has reached the "error count specified value level 1" and copying to the spare disk drive is underway, a parameter "blocked" indicating that copying to a spare disk drive is 20 finished, and a parameter "restoring" indicating that a process of restoring data from another disk drive constituting the disk array (RAID) group to the spare 25 disk drive is underway.

Fig. 13 shows the disk drive management section 230 according to the fourth embodiment and has a copy section 238 in place of the mirror section 235 in Fig.

3. The error monitor section 231 monitors the status of occurrence of errors of a disk drive, instructs initiation of copying to a spare disk drive from a target disk drive when the number of errors occurred in 5 the target disk drive exceeds the "error count specified value level 1", and sets the "blocked" status after copying is done. A blockade/shift section 237 re-sets the value of the "error count specified value level 1".

10 The above is the description of the system structure of the embodiment.

(2) 2 Disk Drives Failure Preventing Operation

This embodiment reduces the probability of occurrence of 2 disk drives failure by dynamically 15 changing the error count specified value which triggers initiation of preventive copying to a spare disk drive in a state where a failure potential of plural disk drives constituting the disk array (RAID) group is high.

The 2 disk drives failure preventing operation is 20 discussed next using a flowchart in Fig. 14.

It is premised on that the error occurrence statuses of the individual disk drives 301 to 307 are counted by the error counter section 232 and are continuously set in the disk drive management table 240.

25 The flowchart in Fig. 11 should be executed independently for the disk drives 301 to 305 constituting a disk array group. It is assumed that the numbers of errors in the disk drives with the "disk

drive Nos." 1 and 3 constituting the disk array (RAID) group are increasing and possible occurrence of a 2 disk drives failure in the disk drives is potentially high.

5 First, the error monitor section 231 determines whether or not the value of the "error counter" in the disk apparatus management table 240 of a disk apparatus to be monitored has reached the "error count specified value level 1" (step 3001). When the "error counter" value has reached the "error count specified value level 1", a process of copying the contents of the disk drive to the spare disk drive and shifting the processing is performed (step 3002). When the "error counter" value has not reached the "error count specified value level 1", it is determined whether or not the "error counter" value has reached the "error count sub specified value" (step 3004). When the "error counter" value has not reached the "error count sub specified value", step 3001 is repeated. When the "error counter" value has reached the "error count sub specified value", it is determined whether or not there is any of those disk drives, excluding the target disk drive, which constitute the disk array (RAID) group and whose error counter value has reached the "error count sub specified value" (step 3005). When there is no such a disk drive, step 3001 is repeated. When there is a disk drive whose error counter value has reached the "error count sub specified value", the values of

the "error count specified value level 1" of all the disk drives constituting the disk array (RAID) group are decreased (step 3006) after which step 3001 is repeated.

5 The re-setting of the value of the "error count specified value level 1" is performed by the blockade/shift section 237. The value to be re-set can be any value, such as an intermediate value between the "error count specified value level 1" and the "error 10 count sub specified value". Although the criterion for the decision in steps 3004 and 3005 is the determination of whether or not there is any of those disk drives, excluding the target disk drive, which constitute the disk array (RAID) group and whose error 15 counter value has reached the "error count sub specified value", it may be the total value of the "error counter" values of all the disk drives constituting the disk array (RAID) group.

20 The 2 disk drives failure preventing operation is carried out as described above.

(3) Advantages

25 The fourth embodiment can provide a highly reliable array-type disk apparatus which copies data or so to a spare disk drive for a possible failure and reduces the probability of occurrence of 2 disk drives failure when a failure potential of plural disk drives constituting the disk array (RAID) group is high.

 Note that the fourth embodiment dynamically

changes the specified value which is the decision reference for the number of errors occurred and may be combined with the first to third embodiments taken singularly or in combination.

5 Further, adapting the data restoring section 270 of the third embodiment in the first and second embodiments can cope with a sector read failure in one disk drive during data restoration which is triggered by a disk drive failure.

10 (Fifth Embodiment)

The fifth embodiment is illustrated below. The fifth embodiment designed to achieve the fourth object of the invention.

15 Fig. 15 is an explanatory diagram showing the structure of storage system, that is, an array-type disk apparatus according to the fifth embodiment of the invention. The array-type disk apparatus of this embodiment comprises a single channel controller or plural channel controllers 1101 each of which has a plurality of host I/Fs for exchanging commands and data with the host computer 100, a cache memory 1301 which temporarily stores input/output data to or from the host computer 100, disk drives 1601 to 1605 to which store input/output data to or from the host computer
20 100, a single disk controller or plural disk controllers A (1401) each having a single or plural disk drive I/Fs 1551, a single disk controller or plural disk controllers B (1402) each likewise having a

single or plural disk drive I/Fs 1552, a shared memory 1302 which can be accessed by both the disk controller A (1401) and disk controller B (1402), and system buses 1201 and 1202 for data transfer and communication among 5 the channel controller 1101, the cache memory 1301, the shared memory 1302 and the disk controllers A (1401) and B (1402). The disk drives D1 (1601), D2 (1602), D3 (1603) and P (1604) have redundancy because of their disk array (RAID) structure.

10 The channel controller 1101 which has received write data from the host computer 100 saves the write data in the cache memory 1301 and instructs the disk controller A (1401) or the disk controller B (1402) to write the write data, located in the cache memory 1301, 15 into the disk drives 1601 to 1604. The channel controller 1101 which has received a data read request from the host computer 100 instructs the disk controller A (1401) or the disk controller B (1402) to read data the disk drives 1601 to 1604 and transfer the 20 data to the cache memory 1301. Having received the instruction, the disk controller A (1401) or the disk controller B (1402) reads data the disk drives 1601 to 1604, transfers the data to the cache memory 1301, then informs the channel controller 1101 of the end of data 25 reading. The informed channel controller 1101 transfers the data from the cache memory 1301 to the host computer 100.

Fig. 16 is a diagram for explaining data

restoration according to the invention, which prevents the occurrence of a 2 disk drives failure, in case where a read error has occurred.

The disk controller A (1401) or disk controller B (1402), which has detected a read error of data D_{1D1} (2001) on the disk drive D1 (1601) updates disk drive information 2101 on the shared memory 1302, reads data D_{2D1} (2002) in the disk drive D2 (1602), data D_{3D1} (2003) in the disk drive D3 (1603), data D_{3D1} (2003) in 10 the disk drive D3 (1603) and data P_{D1} (2004) in the disk drive P (1604) based on the redundant data of data in the disk drive D1 (1601) which has caused a read error, transfers those data to the cache memory 1301 as data D_{2D1} (2302), data D_{3D1} (2303) and data P_{D1} (2304), 15 then restores the data D_{1D1} (2301) in the disk drive D1 (1601) through redundancy calculation using the data D_{2D1} (2302), data D_{3D1} (2303) and data P_{D1} (2304), and stores the restored data D_{1D1} (2301) in the cache memory 1301.

Fig. 17 is a schematic diagram showing the structural elements of the disk drive information (2101).

The disk drive information 2101 comprises a failure counter (error counter) 3001 indicating the 25 number of read errors occurred, a copy counter 3002 indicating the position at which copying to shift data to the disk drive S (1605) is completed, and a disk drive status 3003 indicating information on whether or

not the disk drive is readable/writable. The initial values of the failure counter (error counter) 3001 and the copy counter 3002 are 0, and the initial value of the disk drive status 3003 is the "normal state".

5 Fig. 18 is a flowchart illustrating a status changing process in case where a data read error occurs in the disk drive D1 (1601) while the disk drive status in the disk drive information 2101 is the "normal state".

10 When data reading from the disk drive D1 (1601) is in error, the disk controller A (1401) or the disk controller B (1402) increments the failure counter 3001 in the disk drive information 2101 which concerns the disk drive D1 (1601) in the shared memory 1302 as mentioned above in step 4001. In the next step 4002, it is determined whether or not the failure counter 3001 exceeds a threshold N1. If the failure counter 3001 exceeds the threshold N1, the disk controller A (1401) or the disk controller B (1402) considers that 15 the disk drive D1 (1601) is likely to become completely unreadable in near future, changes the disk drive status 3003 in the disk drive information 2101 to "data being shifted" in step 4003, reads data $D1_{D1}$ (2001) to $D1_{Dm}$ (200n) in the disk drive D1 (1601) onto the cache 20 memory 1301 as data $D1_{D1}$ (2301) to $D1_{Dm}$ (230n) and sequentially writes them in the disk drive S (1605) to thereby shift the data in the disk drive D1 (1601) to the disk drive S (1605) in step 4004. At this time, 25

the copy counter in the disk drive information 2101 is updated to D_m every shifting of data D_{1D_m} ($0 \leq D_m \leq D_n$) to the disk drive S (1605).

5 Fig. 19 is a flowchart illustrating a status changing process in case where a data read error occurs in the disk drive D1 (1601) while the disk drive status in the disk drive information 2101 is "data being shifted".

10 When data reading from the disk drive D1 (1601) is in error, the disk controller A (1401) or the disk controller B (1402) increments the failure counter (error counter) 3001 in the disk drive information 2101 which concerns the disk drive D1 (1601) in the shared memory 1302 as mentioned above in step 5001. In the 15 next step 5002, it is determined whether or not the failure counter (error counter) 3001 exceeds a threshold N2. If the failure counter 3001 exceeds the threshold N2, the disk drive status is changed to "warning", and changes the scheme of reading the data 20 D_{1D_1} (2001) to D_{1D_m} (200n) of the data-shifting disk drive from the disk drive D1 (1601) to the scheme of reading the data from the disk drives D2 to P (1602 to 1604) using the RAID function of the disk array and acquiring restored data through redundancy calculation 25 in step 5004.

Fig. 20 is a flowchart illustrating the scheme of reading data D_{1D_m} ($0 \leq D_m \leq D_n$) from the disk drive D1 (1601) when the disk drive status 3003 in the disk

drive information 2101 is "normal state" or "data being shifted".

In step 6001, data D_{1D_m} is read from the disk drive D1 (1601) and is transferred to the cache memory 1301. In step 6002, it is determined whether a read error has occurred or not. When a read error has occurred, the data D_{1D_m} in the disk drive D1 (1601) is generated using the disk drive D2 (1602), the disk drive D3 (1603) and the disk drive P (1604) which 5 constitute the disk array group having the aforementioned redundancy in step 6003.

The following discusses the scheme of writing data D_{1D_m} ($0 \leq D_m \leq D_n$) in the disk drive D1 (1601) when the disk drive status 3003 in the disk drive 15 information 2101 is "normal state" or "data being shifted". In case where update write data is D_{1D_1} (2301), the disk controller A (1401) or the disk controller B (1402) reads the data D_{1D_1} (2001), located at the associated block position in the disk drive D1 (1601), and stores it on the cache memory 1301 as old data O_{1D_1} (2311). Next, the disk controller A (1401) or the disk controller B (1402) reads the data P_{D_1} (2004) from the disk drive P (1604), and stores it on the cache memory 1301 as old parity data P_{OD_1} (2314). 20 Then, the disk controller A (1401) or the disk controller B (1402) generates new parity data P_{D_1} (2304) through an exclusive-OR operation using the update data D_{1D_1} (2301), the old data O_{1D_1} (2311) and 25

the old parity data P_{D1} (2314), and stores the new parity data P_{D1} (2304) in the cache memory 1301. Next, the disk controller A (1401) or the disk controller B (1402) writes the update data $D1_{D1}$ (2301) in the disk drive D1 (1601) and the disk drive S (1605) and writes the previously generated new parity data P_{D1} (2304).

5 Fig. 21 is a flowchart illustrating the scheme of reading data $D1_{Dm}$ ($0 \leq Dm \leq Dn$) from the disk drive D1 (1601) when the disk drive status 3003 in the disk drive information 2101 is "warning".

10 In step 7001, the data $D1_{Dm}$ in the disk drive D1 (1601) is generated using the disk drive D2 (1602), the disk drive D3 (1603) and the disk drive P (1604) which constitute the disk array group having the aforementioned redundancy in step 6003. When a read error of data DX_{Dm} ($DX: D2$ or $D3$ or P) occurs in any of the disk drives D2 to P (1602 to 1604) in step 7002, the data DX_{Dm} is compared with the copy counter 3002 in the disk drive information 2101 in step 7003. When the data DX_{Dm} is smaller than the copy counter 3002, which means that shifting of this data to the disk drive S (1605) has already been completed, the data $D1_{Dm}$ is read from the disk drive S in step 7004. When the data DX_{Dm} is greater than the copy counter 3002, the data $D1_{Dm}$ is read from the disk drive D1 (1601) in step 7005. At this time, the data DX_{Dm} which had a read error may be restored using the data $D1_{Dm}$. In case where data $D2_{Dm}$ has a read error, $D2_{Dm}$ may be restored using $D1_{Dm}$.

D_{3m} and P_{Dm} constituting the redundant disk array group, a switching medium area may be set in the disk drive D2 (1602) and D_{2m} may be written in that area.

The following discusses the scheme of writing data D_{1m} ($0 \leq D_m \leq D_n$) in the disk drive D1 (1601) when the disk drive status 3003 in the disk drive information 2101 is "warning". In case where write data is D_{1D1} (2301), the disk controller A (1401) or the disk controller B (1402) reads the data D_{2D1} (2002) and data D_{3D1} (2003), which have redundancy at the associated blocks in the disk drive D2 (1602) and the disk drive D3 (1603), and stores them on the cache memory 1301 as old data O_{2D1} (2312) and old data O_{3D1} (2313), respectively. Then, the disk controller A (1401) or the disk controller B (1402) generates new parity data P_{D1} (2304) through an exclusive-OR operation using the update data D_{1D1} (2301), the old data O_{2D1} (2312) and the old data O_{3D1} (2313), and stores the new parity data P_{D1} (2304) in the cache memory 1301. Next, the disk controller A (1401) or the disk controller B (1402) writes the update data D_{1D1} (2301) in the disk drive D1 (1601) and the disk drive S (1605) and writes the previously generated new parity data P_{D1} (2304).

25